arXiv:1510.03203v2 [stat.ML] 14 Oct 2015

VB calibration to improve the interface between phone recognizer and i-vector extractor

Niko Brümmer, AGNITIO Research South Africa

October 15, 2015

1 Introduction

The EM training algorithm of the *classical i-vector extractor* [1, 2] is often incorrectly described as a maximum-likelihood method. The i-vector model is however intractable—the likelihood itself and the hidden-variable posteriors needed for the EM algorithm cannot be computed in closed form. We show here that the classical i-vector extractor recipe is actually a *mean-field variational Bayes (VB)* solution.

This theoretical VB interpretation turns out to be of further use, because it also offers an interpretation of the newer *phonetic i-vector extractor* [4, 5], thereby unifying the two flavours of extractor.

More importantly, the VB interpretation is also practically useful—it suggests ways of modifying existing i-vector extractors to make them more accurate. In particular, in existing methods, the approximate VB posterior for the GMM states is fixed, while only the parameters of the generative model are adapted. Here we explore the possibility of also mildly adjusting (calibrating) those posteriors, so that they better fit the generative model.

In what follows, we introduce notation by summarizing the i-vector model. Then we interpret the classical recipe to deal with this intractable model as a mean-field VB algorithm. Next, we do the same with the newer phonetic i-vector recipe. Finally, we extend the phonetic recipe by allowing calibration of the phone posteriors.

2 The i-vector model

In what follows, speech segments, indexed by s, are represented as variable-length sequences of feature vectors, denoted $\mathbf{\Phi}_s = \{\boldsymbol{\phi}_{st}\}_{t=1}^{T_s}$, where $\boldsymbol{\phi}_{st} \in \mathbf{R}^D$ is a feature vector and T_s is the sequence length. The i-vector model is a generative model for these sequences.

According to the i-vector model, every speech segment is generated by a *different* N-component GMM. The GMM for segment $\mathbf{\Phi}_s$ is parametrized by a set of N weights, covariances and means, in the form:

$$\text{gmm}(\mathbf{x}_s) = \{w_i, \mathbf{C}_i, \boldsymbol{\mu}_i + \mathbf{T}_i\mathbf{x}_s\}_{i=1}^{N} \tag{1}$$

The weights, w_i, and covariances, $\mathbf{C}_i$, remain fixed, while the means vary in a subspace formed by the factor loading matrices, $\mathbf{T}_i$. The result is that every GMM is parametrized by the M-dimensional *i-vector*, $\mathbf{x}_s$. The special GMM, $\text{gmm}(\mathbf{x} = \mathbf{0})$, has means $\boldsymbol{\mu}_i$ and is referred to as the *UBM*.

The i-vector model is completed by specifying a standard normal prior on the i-vectors. The complete i-vector model is parametrized by:

$$\begin{aligned} \Lambda &= \{w_i, \mathbf{C}_i, \boldsymbol{\mu}_i, \mathbf{T}_i\}_{i=1}^{N} \\ &= (\mathbf{U}, \mathbf{T}) \end{aligned} \tag{2}$$

where we have grouped the UBM parameters as $\mathbf{U} = \{w_i, \mathbf{C}_i, \boldsymbol{\mu}_i\}_{i=1}^{N}$ and the factor loading matrices as $\mathbf{T} = \{\mathbf{T}_i\}_{i=1}^{N}$. The i-vector model is shown in graphical model notation in figure 1.

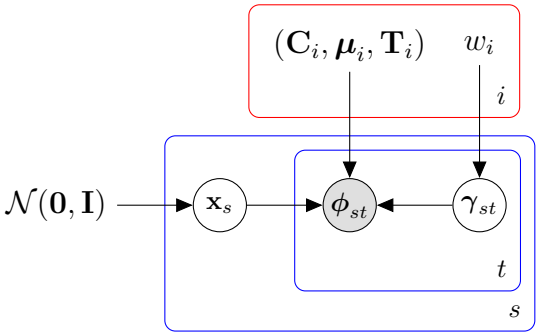

Figure 1: **The i-vector model.** The hidden i-vector is $\mathbf{x}_s$, the hidden GMM state is $\boldsymbol{\gamma}_{st}$ and the observed data point is $\boldsymbol{\phi}_{st}$. The model parameters are in the red plate.

The marginal likelihood, $P(\mathbf{\Phi}_s|\Lambda)$ is intractable, because when only $\mathbf{\Phi}_s$ is given, we cannot jointly sum over the GMM states and integrate out the unknown i-vector. (If the i-vector were also given, the hidden GMM states

would factor over frames and could be summed out in closed form. Or if the state path were given, the i-vector could be integrated out in closed form.)

3 From mean-field VB to the classical i-vector extractor

For the observed sequence $\mathbf{\Phi}_s = \{\boldsymbol{\phi}_{st}\}_{t=1}^{T_s}$, let the hidden GMM state for frame t be represented by a one-hot vector $\boldsymbol{\gamma}_{st}$ and let the whole *path* be denoted $\mathbf{\Gamma}_s = \{\boldsymbol{\gamma}_{st}\}_{t=1}^{T_s}$. The hidden *i-vector* is $\mathbf{x}_s \in \mathbf{R}^M$.

In the true joint hidden variable posterior, $P(\mathbf{x}, \mathbf{\Gamma} | \mathbf{\Phi}_s, \Lambda)$, the i-vector and path are *not* independent, because of explaining away. We obtain the *mean-field VB* approximation [3] by nevertheless forcing such independence in the approximate posterior:

$$P(\mathbf{x}, \mathbf{\Gamma} \mid \mathbf{\Phi}_s, \Lambda) \approx Q_s(\mathbf{\Gamma}, \mathbf{x}) = Q_s(\mathbf{\Gamma}) Q_s(\mathbf{x}) \tag{3}$$

where the factors $Q_s(\mathbf{\Gamma})$ and $Q_s(\mathbf{x})$ are properly normalized distributions that can be found via a variational optimization procedure. More specifically, these factors are found by variational maximization of a functional, known as the *VB lower bound*:

$$\mathcal{L}_s = \left\langle \log \frac{P(\mathbf{\Phi}_s, \mathbf{x}, \mathbf{\Gamma} \mid \Lambda)}{Q_s(\mathbf{\Gamma}) Q_s(\mathbf{x})} \right\rangle_{Q_s(\mathbf{\Gamma}) Q_s(\mathbf{x})} \tag{4}$$

where our notation indicates expectation w.r.t. the variational posterior. For such a mean-field factorization, the optimal factors are known[1] to satisfy:

$$\log Q_s(\mathbf{\Gamma}) = \int Q_s(\mathbf{x}) \log P(\mathbf{\Phi}_s, \mathbf{\Gamma}, \mathbf{x} \mid \Lambda) \, d\mathbf{x} + \texttt{const} \tag{5}$$

$$\log Q_s(\mathbf{x}) = \sum_{\mathbf{\Gamma}} Q_s(\mathbf{\Gamma}) \log P(\mathbf{\Phi}_s, \mathbf{\Gamma}, \mathbf{x} \mid \Lambda) + \texttt{const} \tag{6}$$

where `const` is a place-holder for any extra terms not dependent on the variable of interest (respectively $\mathbf{\Gamma}$ and $\mathbf{x}$). Since these forms are interdependent, this does not solve the optimal factors in closed form—for that you have to iteratively apply the above equations, effectively doing a form of 'co-ordinate ascent'.

Fortunately, by examining these equations in more detail, we can find the functional forms of the optimal factors. We start by expanding (5) and

[1] See e.g. [3], chapter 10.

absorbing the expected value of $\log P(\mathbf{x})$ into `const`, giving:

$$\begin{aligned}\log Q_s(\mathbf{\Gamma}) &= \int Q_s(\mathbf{x}) \log P(\mathbf{\Phi}_s, \mathbf{\Gamma} \mid \mathbf{x}, \Lambda)\, d\mathbf{x} + \texttt{const} \\ &= \int Q_s(\mathbf{x}) \sum_{t=1}^{T_s} \log P(\boldsymbol{\phi}_{st}, \boldsymbol{\gamma}_{st}, \mid \mathbf{x}, \Lambda)\, d\mathbf{x} + \texttt{const} \\ &= \sum_{t=1}^{T} \int Q_s(\mathbf{x}) \log P(\boldsymbol{\phi}_{st}, \boldsymbol{\gamma}_{st}, \mid \mathbf{x}, \Lambda)\, d\mathbf{x} + \texttt{const}\end{aligned} \tag{7}$$

This is enough detail[2] to show that $Q_s(\mathbf{\Gamma})$ factorizes over t, irrespective of the form of $Q_s(\mathbf{x})$. Note that we did not enforce this factorization over t to start with—the original factorization, $Q_s(\mathbf{\Gamma})Q_s(\mathbf{x})$ is sufficient to induce this further factorization.

Indexing the components of the one-hot vectors, $\boldsymbol{\gamma}_{st}$, by the superscripts i, we can expand this distribution as:

$$\log Q_s(\mathbf{\Gamma}_s) = \sum_{t=1}^{T} \log Q_{st}(\boldsymbol{\gamma}_{st}) = \sum_{t=1}^{T} \sum_{i=1}^{N} \gamma_{st}^i \log q_{st}^i \tag{8}$$

where we have introduced the variational parameters, q_{st}^i, which are probabilities that sum to unity over i. Following Bishop, we refer to the q_{st}^i as *responsibilities* [3].

3.1 The VB lower bound

We now expand the lower bound to the form that we shall use for everything that follows. Employing (8) in (4), the *VB lower bound* expands to:

$$\mathcal{L}_s = \left\langle \log \frac{\mathcal{N}(\mathbf{x} \mid \mathbf{0}, \mathbf{I})}{Q_s(\mathbf{x})} + \sum_{t=1}^{T_s} \sum_{i=1}^{N} q_{st}^i \log \frac{w_i \mathcal{N}(\boldsymbol{\phi}_{st} \mid \boldsymbol{\mu}_i + \mathbf{T}_i\mathbf{x}, \mathbf{C}_i)}{q_{st}^i} \right\rangle_{Q_s(\mathbf{x})} \tag{9}$$

This form can be used:

- to derive the formula for $Q_s(\mathbf{x})$;.
- to derive the training algorithm for the extractor parameters, $\mathbf{T}_i$, and also for the UBM parameters if we should need to update them;
- as the objective function for finding the parameters involved in recalibrating the responsibilities.

[2] Later in this document, we shall expand this distribution in more detail.

Finally, we shall also make use of a *simplified lower bound*, obtained from (9) by setting the $\mathbf{T}_i = \mathbf{0}$, which makes $\mathbf{x}$ and $Q_s(\mathbf{x})$ irrelevant. This effectively simplifies the i-vector model to a plain GMM and the resulting lower bound is the same as the well-known EM lower bound for training a GMM:

$$\mathcal{L}_s^{\mathbf{0}} = \sum_{t=1}^{T_s} \sum_{i=1}^{N} q_{st}^i \log \frac{w_i \mathcal{N}(\boldsymbol{\phi}_{st} \mid \boldsymbol{\mu}_i, \mathbf{C}_i)}{q_{st}^i} \tag{10}$$

3.2 The i-vector posterior

Given the responsibilities, q_{st}^i, we can use (8) in (6) to derive the optimal approximate i-vector posterior:

$$\begin{aligned}
\log Q_s(\mathbf{x}) &= \sum_{\boldsymbol{\Gamma}} Q_s(\boldsymbol{\Gamma}) \log P(\boldsymbol{\Phi}_s, \boldsymbol{\Gamma}, \mathbf{x} \mid \Lambda) + \texttt{const} \\
&= \log \mathcal{N}(\mathbf{x} \mid \mathbf{0}, \mathbf{I}) + \sum_{t=1}^{T_s} \sum_{i=1}^{N} q_{st}^i \log \mathcal{N}(\boldsymbol{\phi}_{st} \mid \boldsymbol{\mu}_i + \mathbf{T}_i \mathbf{x}, \mathbf{C}_i) + \texttt{const} \\
&= \mathbf{x}' \left[\sum_{t,i} q_{st}^i \mathbf{T}_i' \mathbf{C}_i^{-1} (\boldsymbol{\phi}_{st} - \boldsymbol{\mu}_i) \right] - \frac{1}{2} \mathbf{x}' \left[\mathbf{I} + \sum_{t,i} q_{st}^i \mathbf{T}_i' \mathbf{C}_i^{-1} \mathbf{T}_i \right] \mathbf{x} + \texttt{const}
\end{aligned} \tag{11}$$

This is a multivariate Gaussian, identical in form to the well-known classical i-vector posterior [2]. The factors in square brackets are the natural parameters of the Gaussian: the *natural mean* (precision times mean); and the *precision* (inverse covariance).

Although we have derived $Q_s(\mathbf{x})$ from (6), the same result could be obtained by working directly from the lower bound (9). To do this, (9) can be interpreted (modulo some constants irrelevant to the optimization) as a negative KL divergence from $Q_s(\mathbf{x})$ to the solution (11). Adopting that solution minimizes the KL divergence and thereby maximizes the lower bound.

3.3 The i-vector extractor

As shown above, given the model parameters, $\Lambda = (\mathbf{U}, \mathbf{T})$, the *full* mean-field VB recipe for inferring the i-vector would iteratively update the q_{st}^i and $Q_s(\mathbf{x})$, using (5) and (11), until convergence.

The *classical i-vector extractor algorithm* approximates the full recipe with a non-iterative shortcut:

1. *Fix* the responsibilities by equating them to the UBM state posteriors. This is the same as optimizing $\mathcal{L}_s^{\mathbf{0}}$ w.r.t. the q_{st}^i and can be done in one closed-form 'E-step'. This step is usually called the *UBM alignment*.

2. Compute $Q_s(\mathbf{x})$, using (11).

3.4 The extractor training algorithm

To train the model parameters, Λ, using the *full mean-field VBEM*, would involve iterative application of three steps:[3]

- Update the q_{st}^i, using (5).
- Update the $Q_s(\mathbf{x})$, using (11).
- Update Λ, as the value that maximizes $\sum_s \mathcal{L}_s$.

In this mean-field VBEM algorithm, the lower bound and all the updates can be computed in closed form.

The *classical 'EM-algorithm'* for training the extractor parameters simplifies the full VBEM algorithm by an approximating short-cut:

1. *Fix* $\mathbf{U}$ by training an ordinary GMM on the training data, without considering any i-vector mechanism.
2. *Fix* the q_{st}^i, using the same UBM alignment as above.
3. After some appropriate initialization, iterate:
   - Update the $Q_s(\mathbf{x})$, using (11)
   - Update $\mathbf{T}$, by maximizing $\sum_s \mathcal{L}_s$.

4 The phonetic i-vector extractor as a VB algorithm

In the classical recipe above, we (i) fixed the UBM, which then (ii) supplied the responsibilities, which allowed (iii) training $\mathbf{T}$. Here, we swap the order of steps (i) and (ii), by first fixing the responsibilities via some other means, then finding the UBM parameters and then $\mathbf{T}$.

The fact that the data is observed, places strong constraints on probability distributions for the data. Things like the dimensionality of the data are fixed by the observation. With the hidden variables, we have much more freedom. A wide range of different shapes and sizes of hidden variables can all provide reasonably good explanations for the observed data. Within limits,

[3] not necessarily in that order

we could explore different hidden-variable dimensionalities, with different kinds of support (e.g. discrete vs continuous).

In the case of a the *phonetic i-vector extractor*, we keep the state discrete, but we usually increase N somewhat. We also force the hidden states to correspond to triphone classes (senones). In short, we run a senone classifier on the *raw* speech segment (not on Φ_s, the phone recognizer computes its own features) and the result is a factorial posterior distribution[4] over senone classes for every frame, with the frame rate corresponding to that of the original i-vector extractor.

In the case of text-independent speaker (or language) recognition, the phone recognizer could be as complex as a full LVCSR system, including a language model; and for text-dependent speaker recognition, it could use forced alignment, using the known text. Or, it could simply be a raw phone recognizer. These details are not important for this analysis, as long as the resulting phone (senone) posterior is factorial over frames.

The phonetic i-vector extractor recipe can be summarized as follows:

1. *Fix* the responsibilities, q_{st}^i, by running a senone recognizer on the speech.

2. *Fix* the UBM parameters, $\mathbf{U} = \{w_i, \boldsymbol{\mu}_i, \mathbf{C}_i\}_{i=1}^N$ by maximizing $\sum_s \mathcal{L}_s^{\mathbf{0}}$. This can be done in closed form in a single 'M-step'.[5]

3. As above: initialize, then iterate:

   - Update the $Q_s(\mathbf{x})$, using (11)
   - Update $\mathbf{T}$, by maximizing the full objective, $\sum_s \mathcal{L}_s$.

In the above recipe, $\mathbf{U}$ is fixed. With a bigger computational budget, it is however possible to maximize $\sum_s \mathcal{L}_s$ jointly w.r.t. $\mathbf{U}$ and $\mathbf{T}$ with an iterative EM-algorithm, with closed-form updates. In this EM-algorithm, the E-step updates the i-vector posterior in closed form, while the M-step updates both $\mathbf{T}$ and $\mathbf{U}$ in closed form. This may be an interesting experiment to perform—we could potentially obtain a more accurate model, with the caveat that the risk of overfitting increases.

[4] With wide input contexts and phonetic and language models, the senone posteriors should not be factorial, but current technology does produce factorial senone posteriors, relying on later processing to link them.

[5] Although this also updates the w_i, we don't need them for this recipe. We shall however need them for the next recipe when we recalibrate the responsibilities.

5 Recalibrating the phone posteriors with VB

Now we employ the VB interpretation to generalize the phonetic i-vector extractor training, with the purpose of improving the 'collaboration' between the phone recognizer and the rest of the extractor. In both of the above VB algorithms (classical and phonetic), the approximate posterior was fixed before finalizing all of the extractor parameters. However, the general VB framework allows for joint optimization of approximate posterior and model parameters.

In the case of the phone recognizer, we don't want to make big changes to the phone recognizer (presumably it is already working well, and it might be a third party product which is difficult to adjust). But, we can allow calibration of the phone posteriors. This is relatively easy to do and preserves the information contained in the original posteriors. By optimizing the calibration via the VB lower bound, the KL divergence between approximate and actual posteriors will be reduced, allowing them to work better together.

5.1 The calibration transformation

In this case, we denote the *fixed* posteriors supplied by the phone recognizer as, $\tilde{q}^i_{st}$. Now we allow our VB responsibilities, q^i_{st}, to be *adjustable*, but only mildly so. That is, we constrain the responsibilities by deriving them from the fixed phone posteriors via the *calibration transformation*:

$$\{q^i_{st}\}_{i=1}^N = \text{softmax}\left(\{\alpha \log \tilde{q}^i_{st} + \beta_i\}_{i=1}^N\right) \tag{12}$$

where only the *calibration parameters*, $C = (\alpha, \{\beta_i\}_{i=1}^N)$ are adjustable. The softmax function renormalizes the distribution after the adjustment in the log domain. In the VB context, C may be referred to as a *variational* parameter. The scale factor α adjusts the sharpness of the distribution, while the offsets β_i play the same role as adjusting the prior.

5.2 The calibration objective

Here we want to update the responsibilities, q^i_{st}, by adjusting the calibration parameters, when the model parameters, $\Lambda = (\mathbf{U}, \mathbf{T})$, and the $Q_s(\mathbf{x})$ are given. To do this, we conveniently massage the lower bound (9) into the form of a KL divergence. It is understood that the maximization below is w.r.t. the q^i_{st} only and that the maximand may be conveniently adapted from

line to line by adding or subtracting terms independent of q^i_{st}:

$$\begin{aligned}
&\max \sum_s \mathcal{L}_s \\
&= \max \sum_s \sum_{t=1}^{T_s} \sum_{i=1}^{N} q^i_{st} \left\langle \log \frac{w_i \mathcal{N}(\boldsymbol{\phi}_{st} \mid \boldsymbol{\mu}_i + \mathbf{T}_i \mathbf{x}, \mathbf{C}_i)}{q^i_{st}} \right\rangle_{Q_s(\mathbf{x})} \\
&= \max \sum_s \sum_{t=1}^{T_s} \sum_{i=1}^{N} q^i_{st} \log \frac{w_i \ell^i_{st}}{q^i_{st}} \\
&= \max \sum_s \sum_{t=1}^{T_s} \sum_{i=1}^{N} q^i_{st} \log \frac{r^i_{st}}{q^i_{st}}
\end{aligned} \tag{13}$$

where we have introduced the likelihoods,

$$\ell^i_{st} = w_i \exp \left[\langle \log \mathcal{N}(\boldsymbol{\phi}_{st} \mid \boldsymbol{\mu}_i + \mathbf{T}_i \mathbf{x}, \mathbf{C}_i) \rangle_{Q_s(\mathbf{x})} \right] \tag{14}$$

which when normalized give the *optimal responsibilities*:

$$r^i_{st} = \frac{\ell^i_{st}}{\sum_{j=1}^{N} \ell^j_{st}} \tag{15}$$

By this normalization, the objective is now in the form of a sum of negative KL divergences, each of which can be minimized (zeroed) by setting $q^i_{st} = r^i_{st}$. If we were doing the full, unconstrained mean-field VB iteration, this is exactly what we would do. Here however, we want to retain the correspondence between senones and our hidden states, so we place a strong constraint, namely (12), on the responsibilities.

In the practical implementation of this calibration it is good to see that (14) can be computed in closed form, since it is a Gaussian expectation of a quadratic expression. To optimize (13) w.r.t. the calibration parameters C, we will have to resort to general purpose, iterative, numerical optimization, for example BFGS [6].

5.3 The algoritm

We can now assemble the full algorithm:

1. *Fix* the senone posteriors $\tilde{q}^i_{ts}$.
2. Initialize C: $\alpha = 1$ and $\beta_i = 0$. This implies also, $q^i_{st} = \tilde{q}^i_{st}$.
3. Initialize $\mathbf{U}$ as before, by maximizing $\sum_s \mathcal{L}^{\mathbf{0}}_s$.

4. Initialize $\mathbf{T}$ and iterate:
   - Update $Q_s(\mathbf{x})$, using (11).
   - Update C using iterative numerical maximization of (13).
   - Update $\mathbf{T}$ and $\mathbf{U}$ by maximizing $\sum_s \mathcal{L}_s$.

6 Beyond calibration

As noted above, with hidden variables, we have much freedom. We could generalize the calibration recipe above in various ways:

- Generalize α by a diagonal matrix, or a full square matrix, or even a rectangular matrix. (A rectangular matrix can reduce the number of senone classes to a smaller number of hidden variable states.)
- Fuse more than one phone recognizer (e.g. different ones for different languages), by weighting the log posteriors of each with its own rectangular matrix. With rectangular matrices, we can fuse log posteriors of different sizes.
- Smoothing log-posteriors over time, by 'fusing' q_{ti} with posteriors of neighbouring frames.
- We could also go in the other direction and use VB to design lightweight hidden variable posteriors calculators. One way to do this would be to first find the best-performing (presumably computationally expensive) posterior calculator, with its associated model parameters, $\mathbf{U}$ and $\mathbf{T}$. Then fix these model parameters and optimize the (variational) parameters of a computationally cheaper posterior calculator. VB will ensure that the cheaper calculator matches the model parameters as well as possible.

References

[1] Dehak, N., Dehak, R., Kenny, P., Brummer, N., Ouellet, P and Dumouchel, P., "Support Vector Machines versus Fast Scoring in the Low-Dimensional Total Variability Space for Speaker Verification", Interspeech 2009, Brighton, UK, September 2009.

[2] Dehak, N., Kenny, P., Dehak, R., Dumouchel, P and Ouellet, P. "Front-End Factor Analysis for Speaker Verification", IEEE Transactions on Audio, Speech and Language Processing, 19(4), pp. 788-798, May 2011.

[3] Christopher M. Bishop, *Pattern Recognition and Machine Learning*, Springer, 2006.

[4] Y. Lei, N. Scheffer, L. Ferrer, and M. McLaren, "A novel scheme for speaker recognition using a phonetically aware deep neural network", in ICASSP 2014.

[5] Daniel Garcia-Romero and Alan McCree, "Insights into Deep Neural Networks for Speaker Recognition", Interspeech, Dresden 2015.

[6] Jorge Nocedal and Stephen J. Wright, *Numerical Optimization*, 2nd edition, Springer 2006.